



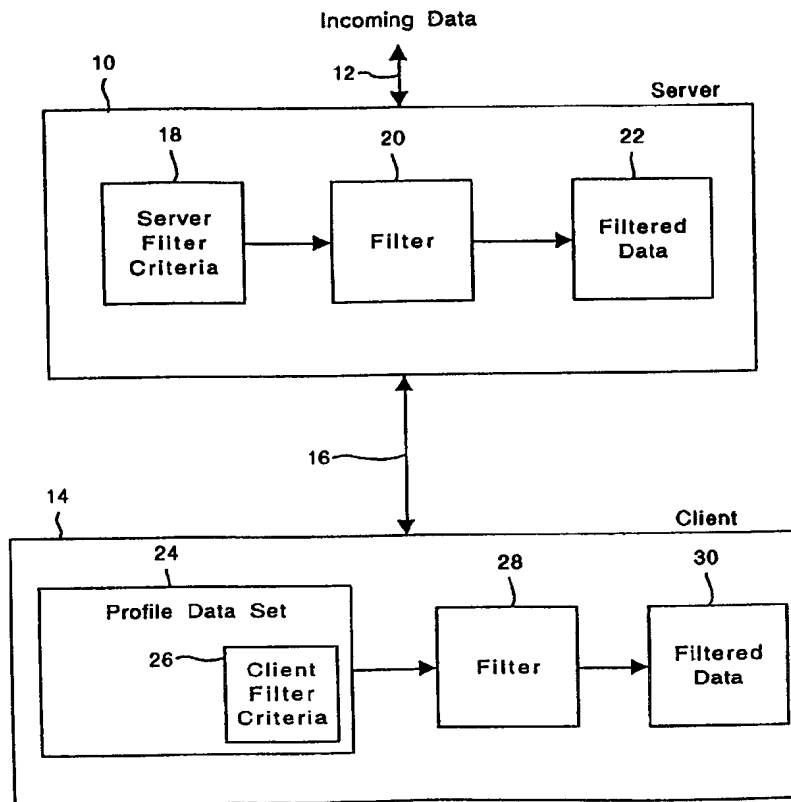
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 17/30, 13/36, 15/00, H04N 1/413		A1	(11) International Publication Number: WO 99/28838 (43) International Publication Date: 10 June 1999 (10.06.99)
(21) International Application Number: PCT/US98/25647 (22) International Filing Date: 2 December 1998 (02.12.98) (30) Priority Data: 08/985,389 4 December 1997 (04.12.97) US 09/198,337 23 November 1998 (23.11.98) US (71) Applicant: AVEO, INC. [US/US]; Suite 208, 2901 Tesman Drive, Santa Clara, CA 95054 (US). (72) Inventor: HOFMANN, William, D.; 1408 Carleton Street, Berkeley, CA 94702 (US). (74) Agents: SPONSELLER, Allan, T. et al.; Blakely, Sokoloff, Taylor & Zafman, 7th floor, 12400 Wilshire Boulevard, Los Angeles, CA 90025-1026 (US).			(81) Designated States: AL, AM, AT, AT (Utility model), AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), EE, EE (Utility model), ES, FI, FI (Utility model), GB, GD, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: METHOD AND APPARATUS FOR MULTI-STAGE DATA FILTERING BY A SINGLE DEVICE

(57) Abstract

A system for filtering data in multiple stages at a single location without exposing private information to untrusted servers includes receiving identifying information corresponding to a first set of data, at a first device (10), from a second device (14). The first set of data (22) is filtered based on both the first filter criteria (18) and an identifier information, in order to identify a first set of filtered data (22). Subsequence to the filtering of the first set of data (22), the first set of filtered data (22) is received from the second device (14). The first set of filtered data (22) is then filtering based on a second filter criteria (24), wherein the filtering of the first set of filtered data (22) generates a second set of filtered data (30), and wherein the second filter criteria (26) is different from the first filter criteria (18).



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

Method and Apparatus for Multi-Stage Data Filtering by a Single DeviceBACKGROUND OF THE INVENTION1. Related Applications

This is a continuation-in-part of Application No. 08/985,389, filed December 4, 1997, entitled "Multi-Stage Data Filtering System".

2. Field of the Invention

The present invention relates to a data filtering system. More specifically, the present invention provides a system capable of filtering data in multiple stages by a single device, with each stage of filtering using different filtering criteria.

3. Background Information

The increased use of networks (such as the Internet) and networking technology has increased the quantity of data presented to individuals and organizations on a day-to-day basis. This data may be in the form of advertisements, news articles, and other information from any number of data sources. Although much of this data may be of interest to particular individuals and organizations, a significant portion of the data is generally of little or no value to the recipient. For example, the data may be related to a subject that is of no interest to the recipient or related to a type of product that the recipient does not use and does not intend to purchase.

Existing systems are available for selecting data to be provided to a particular user based on criteria that is supplied actively or passively by the user. These existing systems perform various filtering operations on a server to select the data to be provided to a particular user. Since these filtering operations are performed on a centralized server, the server must contain the necessary filtering criteria to select the data. These existing systems limit the effectiveness of the filtering operation because certain criteria necessary for proper filtering is confidential or private to the user and is not disclosed to the server. Since the server does not have this private information, it cannot adequately filter out all of

-2-

the irrelevant data. For example, if a user does not indicate their age to the server, then the server cannot filter data that is directed at a particular age group. As a result, the user receives all data regardless of whether the data is relevant to a person in the user's age group.

Since the server is unable to filter data based on private criteria not provided to the server by the user, the user may receive a significant amount of irrelevant data. This irrelevant data is time consuming to review and creates a distraction from the user's normal work or activities. Since many servers that provide data filtering operations may not be trustworthy with respect to private information, many users are unwilling to provide private information to these servers. As a result, the user receives a significant amount of unwanted data.

Other known systems for filtering data provide all data from all sources to the client, which then filters the data based on information provided by the user of the client. This approach significantly increases the amount of data received by the client and increases the bandwidth or transmission time required to transmit the data to the client from the data sources. The increase in data received by the client also increases the local storage requirements.

It is therefore desirable to provide an improved data filtering system capable of filtering out data that is not relevant to a particular user, without compromising the user's privacy.

SUMMARY OF THE INVENTION

The present invention is related to a system for filtering data in multiple stages at a single location without exposing private information to untrusted servers. According to one aspect of the present invention, identifier information corresponding to a first set of data is received, at a first device, from a second device. The first set of data is filtered based on both the first filter criteria and the identifier information, in order to identify a first set of filtered data. Subsequent to the filtering of the first set of data, the first set of filtered data is received from the second device. The first set of filtered data is then filtered based on a second filter criteria, wherein the filtering of the first set of filtered data generates a second set

of filtered data, and wherein the second filter criteria is different from the first filter criteria.

According to another aspect of the present invention, data is filtered based on both a first set of identifiers corresponding to the data and a first filter criteria in order to generate a first set of filtered data identifiers. Additionally, at least a portion of the first set of filtered data identifiers is filtered based on a second filter criteria in order to generate a second set of filtered data identifiers.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example in the following drawings in which like references indicate similar elements. The following drawings disclose various embodiments of the present invention for purposes of illustration only and are not intended to limit the scope of the invention.

Figure 1 illustrates an embodiment of a multi-stage data filtering system.

Figure 2 is a flow diagram illustrating an embodiment of a procedure for performing multi-stage data filtering.

Figure 3 is a flow diagram illustrating another embodiment of a procedure for performing multi-stage data filtering.

Figure 4 illustrates another embodiment of a multi-stage data filtering system.

Figure 5 is a flow diagram illustrating another embodiment of a procedure for performing multi-stage filtering.

Figure 6 illustrates an embodiment of a profile data set for use with the present invention.

Figure 7 illustrates exemplary profile data elements related to user-specific information according to one embodiment of the present invention.

Figures 8A and 8B illustrate exemplary server filter criteria and client filter criteria generated from the profile data elements shown in Figure 7.

Figure 9 illustrates another embodiment of a multi-stage data filtering system.

Figure 10 illustrates another embodiment of a multi-stage data filtering system.

Figure 11 illustrates another embodiment of a multi-stage data filtering system.

Figures 12A, 12B, and 12C illustrate exemplary filter criteria for use in the multi-stage data filtering system shown in Figure 11.

Figure 13 illustrates another embodiment of a multi-stage data filtering system.

Figure 14 illustrates an embodiment of a computer system that can be used with the present invention.

Figure 15 illustrates an embodiment of a computer-readable medium containing various sets of instructions, code sequences, configuration information, and other data used by a computer or other processing device.

DETAILED DESCRIPTION

The following detailed description sets forth numerous specific details to provide a thorough understanding of the invention. However, those of ordinary skill in the art will appreciate that the invention may be practiced without these specific details. In other instances, well-known methods, procedures, protocols, components, and circuits have not been described in detail so as not to obscure the invention.

The present invention is related to a system capable of filtering data for a particular user (also referred to as a data recipient) without compromising that user's privacy. The invention provides a unified data filtering process such that data filtering is performed in multiple stages by a single device, with different filtering criteria used for each stage. In a first stage, data can be filtered by a device using both non-private filtering criteria and identifying information corresponding to the data. The data (or identifying information) can then be provided to the device, which can filter the data based on more private information about the user or organization. Any number of filtering stages may be utilized by a single device. By limiting private filtering criteria to trusted devices, a significant amount of

-5-

unwanted data is eliminated without compromising the user's privacy. Additionally, by reducing the number of devices which perform filtering, the management of filtering devices is more centralized and thus simplified.

Throughout this detailed description of the invention, various embodiments are discussed that include a client coupled to one or more servers. The teachings of the present invention are applicable to any type of device containing a processor or a controller capable of executing instructions. Thus, the clients and servers discussed herein can be any type of computing device, including desktop or laptop computers, personal digital assistants (PDAs), set-top boxes, or devices containing embedded controllers or embedded processors. Further, any type of communication link and communication medium can be used to communicate information between two or more devices.

Particular data filtering procedures are described below that utilize a profile data set to generate filter criteria for servers and clients. However, it will be appreciated that any method or procedure for filtering data can be used with the present invention. Further, any number of filtering parameters or attributes may be used to filter data at any number of data filtering stages. Additionally, the present invention can be used with any type of data (e.g., text, graphics, product updates (such as software updates), or executable instructions) and with data received from any data source or sources.

Figure 1 illustrates an embodiment of a multi-stage data filtering system. A server 10 receives incoming data on a communication link 12. Communication link 12 may be a network communication link or any other link capable of communicating data between two or more devices. Server 10 communicates with a client 14 using a communication link 16. Communication link 16 may be a link through a network or any other link capable of propagating data between server 10 and client 14. Communication links 12 and 16 may use any type of communication medium, such as, but not limited to, wires, fiber optic cables, or wireless communication systems.

Server 10 includes server filter criteria 18, which provides the filtering criteria used by a filter 20 to filter the incoming data. Server 10 may be an

-6-

untrusted server with which users are unwilling to share private information. In this situation, server filter criteria 18 contains public information (i.e., public filtering criteria) that the user is willing to share with the server. Additional details regarding server filter criteria 18 and the operation of filter 20 are provided below. Filter 20 generates filtered data 22 as a result of applying server filter criteria 18 to the incoming data. Filtered data 22 is generally a subset of the incoming data received on communication link 12. However, in certain situations, filtered data 22 is a null set of data if filter 20 removes (i.e., filters out) all of the incoming data. In other situations, all incoming data may pass through filter 20, such that filtered data 22 contains all incoming data. Upon completion of the filtering operation performed by filter 20, filtered data 22 is provided to client 14 using communication link 16.

Client 14 contains a profile data set 24, which includes client filter criteria 26. In this embodiment, client 14 is trusted and, therefore, client filter criteria 26 may include private information that is not shared with server 10. Profile data set 24 contains all profile data associated with a particular user or organization. This profile data is used to generate server filter criteria 18 and client filter criteria 26. In the embodiment shown in Figure 1, profile data set 24 contains server filter criteria 18 and client filter criteria 26. In alternate embodiments, profile data set 24 may include filter criteria associated with a particular class of users or a particular role that a user performs. Additional details regarding profile data sets are provided below with respect to Figures 6-8.

Client 14 also includes a filter 28 that applies client filter criteria 26 to filtered data 22 received from server 10 on communication link 16. Filter 28 generates a set of filtered data 30, representing the incoming data that meets both server filter criteria 18 and client filter criteria 26. Filtered data 30 is then provided to the user of client 14 for viewing or other processing. To maintain the privacy of the information contained in the profile data set, the results of the filtering process at any particular level of trust are not provided to a device or filtering process having a lower level of trust.

-7-

As shown in the filtering system of Figure 1, profile data set 24 is contained in client 14. Thus, only the data that is public (i.e., not confidential or private) is shared with server 10. The remaining filter criteria are stored on the client and is not exposed to or otherwise provided to the server. Thus, the single profile data set 24 provides a unified system for filtering incoming data on both server 10 and client 14.

The embodiment of Figure 1 represents a unified two-stage data filtering system. However, the teachings of the present invention may be applied to a data filtering system having any number of data filtering stages. An example of a unified three-stage data filtering system is illustrated in Figure 11 and discussed below. Additionally, Figure 1 shows a single client 14 coupled to server 10. In other embodiments of the invention, a particular server may be coupled to multiple clients and contain separate filter criteria for each client that receives data from the server.

Figure 2 is a flow diagram illustrating an embodiment of a procedure for performing multi-stage data filtering. The procedure illustrated in Figure 2 may be used, for example, with the data filtering system illustrated in Figure 1. At step 40, a profile data set is generated and stored on a client (e.g., client 14 in Figure 1). Additional details regarding the profile data set are discussed below with reference to Figures 6-8. At step 42, the procedure determines the level of trust associated with a server (e.g., server 10 in Figure 1). For example, a server located inside (i.e., on the corporate side) of a firewall may have a high level of trust and security, but a server located outside the firewall may be untrustworthy and is assigned a low level of trust. The level of trust associated with a particular server determines the type of profile data that is shared with that server for data filtering purposes. If a level of trust is not assigned to a particular server, then the server may be assigned a default level of trust (e.g., an untrusted server).

At step 44 of Figure 2, the client transmits profile data elements associated with the server's level of trust to the server. These profile data elements are referred to as the server filter criteria. The server filter criteria is stored within the server (e.g., in a register or other data storage mechanism). The server filter criteria

-8-

may be stored temporarily or permanently. At step 46, the procedure determines whether incoming data was received by the server. If no data was received, the procedure returns to step 46 to repeatedly test for incoming data. As an alternative to repeated testing for incoming data, the procedure may use a "trigger" that causes the procedure to continue to step 48 when incoming data is detected.

At step 48, the procedure filters the incoming data on the server using the server filter criteria. Step 50 transmits the filtered data, if any, from the server to the client. At step 52, the procedure filters data received by the client using the profile data elements associated with the client. These profile data elements are referred to as the client filter criteria. Finally, step 54 processes the filtered data, if any, generated by the client. This processing may include displaying the data to a user or notifying the user of the received data. If either the filtering performed by the server at step 48 or by the client at step 52 eliminates all data, then the procedure terminates without notifying the user.

Figure 3 is a flow diagram illustrating another embodiment of a procedure for performing multi-stage data filtering. The procedure illustrated in Figure 3 may be used, for example, with the data filtering system illustrated in Figure 1. The procedure of Figure 3 is similar to the procedure discussed above with respect to Figure 2, but transmits profile data elements to the server after the receipt of incoming data instead of prior to the receipt of incoming data. At step 60, a profile data set is generated and stored on the client. Step 62 determines whether incoming data has been received by the server. If incoming data has not been received, then the procedure returns to step 62 to continue testing for incoming data. Alternatively, a "trigger" can be used that causes the procedure to continue to step 64 when incoming data is detected.

When incoming data is received, the procedure continues to step 64, in which the server requests filter criteria from the client. In response to the server's request for filter criteria, the client determines the level of trust associated with the requesting server at step 66. At step 68, the client transmits profile data elements associated with the server's level of trust to the server. These profile data elements are referred to as the server filter criteria. In a particular embodiment of

the invention, the server discards the server filter criteria after filtering the received data. In an alternate embodiment of the invention, the server may store the server filter criteria for use with the next incoming data. In this alternate embodiment, the client may update the server with new server filter criteria each time the server filter criteria changes.

At step 70 of Figure 3, the incoming data is filtered on the server using the server filter criteria. Step 72 transmits the filtered data, if any, from the server to the client. At step 74, the data received by the client is filtered using the profile data elements associated with the client (referred to as the client filter criteria). The filtered data, if any, generated by the client is then processed at step 76. As discussed above, this processing may include displaying the filtered data to the user or notifying the user of the received data. If either the filtering performed by the server at step 70 or by the client at step 74 eliminates all data, then the procedure terminates without notifying the user.

Embodiments of the present invention execute the procedures described above with respect to Figures 2 and 3 continually (e.g., in a background mode). Therefore, the client and server(s) may exchange filter criteria, filtered data, and other information while the client is executing other applications or procedures.

Figure 4 illustrates another embodiment of a multi-stage data filtering system. The system of Figure 4 operates similarly to that of Figure 1, except that the multi-stage filtering is all done on the client 84 rather than being split between the client 84 and the server 80. Server 80 receives incoming data on a communication link 82 and communicates with client 84 using a communication link 86. Server 80 takes identifier information corresponding to the incoming data it receives and forwards the information to client 84. The identifier information identifies (either explicitly or inherently) its corresponding data as well as enough information corresponding to the received data in order to allow client 84 to accurately filter the data, and can vary for different types or pieces of data. Client 84 filters the data using the identifier information and server filter criteria 88 stored locally at client 84 as part of profile data set 94. Filter 98 generates a filtered

-10-

request list as a result of applying server filter criteria 88 to the identifier information.

The filtered request list is returned to data selector 90 of server 80, as an information request targeted at data selector 90. The data corresponding to the identifier information that is on the filtered request list is identified as filtered data 92 and is forwarded by server 80 to client 84 via communication link 86. Upon receipt, filter 98 applies client filter criteria 96 to the filtered data 92. Filter 98 generates a set of filtered data 100, representing the incoming data that meets both server filter criteria 88 and client filter criteria 96. Filtered data 100 is then provided to the user of client 84 for viewing or other processing.

Alternatively, client 84 may perform filtering for multiple stages at the same time. For example, upon receipt of the identifying information from server 80, client 84 may apply both the server filter criteria to generate a first filtered request list and then apply the client filter criteria based on this first list to generate a second filtered request list. Client 84 then forwards the first list to server 80 and keeps the second list locally, so that when filtered data is received from server 80, client 84 already knows which data is to be discarded and which is to be processed. This filtering based on the client filter criteria may be done prior to sending the first list to server 80, or alternatively at other times, such as while the first list is being sent to and the filtered data is being received from server 80.

According to alternate embodiments, client filter criteria 96 includes the server filter criteria 88, resulting in a single criteria set rather than two sets. In such embodiments, the data elements within client filter criteria 96 include indications of whether the filter applies to the server stage or the client stage of filtering (e.g., the privacy characteristics discussed below with reference to Figure 7).

Figure 5 is a flow diagram illustrating another embodiment of a procedure for performing multi-stage filtering. The procedure illustrated in Figure 5 may be used, for example, with the data filtering system illustrated in Figure 4. The procedure of Figure 5 is similar to the procedures discussed above with respect to Figures 2 and 3, except that multiple stages of filtering is done at a single device

-11-

(e.g., the client) rather than at multiple devices (e.g., both the client and the server). At step 110, profile data sets are generated and stored on a client (e.g., client 84 in Figure 4). These profile data sets include both client filter criteria and server filter criteria. Step 112 determines whether incoming data has been received by the server. If incoming data has not been received, then the procedure returns to step 112 to continue testing for incoming data. Alternatively, a "trigger" can be used that causes the procedure to continue to step 114 when incoming data is detected.

When incoming data is received the procedure continues to step 114, in which the server forwards identifier information for the incoming data to the client. At step 116, the client filters the data based on the received identifier information and the server's level of trust, step 116. This filtering is based on the server filter criteria. At step 118, an information request based on the filtered identifier information is transmitted to the server from the client. This information request is a request for all of the data which satisfies the server filter criteria.

At step 120, the server transmits the requested data, if any, to the client. Upon receipt of the data, the client filters the data using profile data elements associated with the client (that is, using the client filter criteria), step 122. The filtered data, if any, generated by the client is then processed at step 124. As discussed above, this processing may include displaying the filtered data to the user or notifying the user of the received data. If the filtering performed by the client at either step 116 or step 122 eliminates all data, then the procedure terminates without notifying the user.

Thus, it can be seen that multiple stages (both the "server stage" and the "client stage") of filtering is done at the client even though not all data is forwarded to the client. This helps protect the user's privacy because no profile information is stored at the server. Furthermore, because the data is still being filtered in multiple stages, the server can only infer information about the user based on the filtering done by the client using the server filter criteria. The server still has no knowledge of which data, if any, is discarded by the client using the client filter criteria.

-12-

Figure 6 illustrates an embodiment of a profile data set 130 for use with the present invention. In one embodiment of the invention, a separate profile data set 130 is provided for each client (or each user). Profile data set 130 includes a set of profile data elements 132 that are related to user-specific information (e.g., age, occupation, or marital status). Profile data set 130 also includes a set of profile data elements 134 that are related to one or more user roles. A user role can be, for example, "professor" or "Vice President of Engineering." Profile data elements 134 related to a user role identify characteristics or attributes associated with that role, rather than an individual person. Therefore, all users performing a particular role may use profile data elements 134 rather than or in addition to entering those attributes along with their user-specific information. Furthermore, the attributes associated with a particular role can be updated once rather than updating each user's specific information. If a particular user performs multiple roles, then that user's profile data set 130 will contain profile data elements related to all of the roles performed by the user.

Profile data set 130 further includes a set of profile data elements 136 that are related to one or more user classes. A user class can be, for example, "marketing" or "engineers." Profile data elements 136 related to a user class identify characteristics or attributes associated with a class of users. Therefore, all users that are members of a particular class can use profile data elements 136 rather than entering those attributes along with their user-specific information. Additionally, the attributes associated with a particular class can be updated once rather than updating each member's specific information. If a particular user is a member of multiple classes, then that user's profile data set 130 will contain profile data elements related to all of the classes of which the user is a member. Additionally, a particular user may override the value associated with an attribute associated with a role or a class. For example, a role "Software Engineering Manager" may have an attribute "job level" with a value "grade 1." If a particular user performing the role of Software Engineering Manager has a job level "grade 2," that user's profile data set will contain an entry for the "job level" – "grade 2" pair that overrides the value provided by the role. Thus, the values associated

-13-

with role or class attributes may operate as default values that can be changed by a user's profile data set.

As shown in Figure 6, profile data elements 134 related to user roles and profile data elements 136 related to user classes are stored within profile data set 130. In alternative embodiments of the invention, a pointer or similar mechanism is provided in profile data set 130 that identifies a centralized storage location for the profile data elements related to user roles or user classes. The use of profile data elements related to user roles and user classes is optional. In alternative embodiments of the invention, profile data set 130 may include only profile data elements 132 related to user-specific information.

Figure 7 illustrates exemplary profile data elements related to user-specific information (e.g., profile data elements 132 in Figure 6). The data elements shown in Figure 7 are arranged into a table format for purposes of illustration. However, the actual data elements may be stored in any configuration using any data structure. The data elements in Figure 7 include several attribute-value pairs (i.e., a value associated with each attribute). Additionally, a privacy characteristic is associated with each attribute-value pair. For example, the attribute "name" has a value "John Doe" and an associated privacy characteristic "Public." Thus, the user's name is John Doe and the user has made their name public. Public attributes are provided to all servers (whether the server is considered trustworthy or untrustworthy). The employer attribute has a value "Acme Corp." and has an associated privacy characteristic "Semi-Private." A "Semi-Private" privacy characteristic indicates that the attribute is only provided to trustworthy servers (i.e., not provided to untrustworthy servers). Trustworthy servers may be those servers located inside a corporate firewall and untrustworthy servers may be those servers located outside the corporate firewall. A third privacy characteristic, "Private," indicates that the attribute is only provided to clients, and is not provided to any server, whether trusted or untrusted. The example of Figure 7 contains three different levels of privacy (Public, Semi-Private, and Private). However, in alternate embodiments of the invention, any number of privacy levels

-14-

may be provided. As discussed in greater detail below, the number of privacy levels does not necessarily equal the number of filtering stages.

Alternatively, in embodiments where multiple stages of filtering are being performed by a single device, the privacy characteristic indicates which filter criteria are made aware of the attribute-value pair regardless of the location of that filter criteria. For example, public attributes are made available to all server filter criteria, regardless of whether the filter criteria is located at (and the filtering performed at) a client or a server. Additionally, semi-private attributes are made available to trusted server filter criteria, regardless of whether the filter criteria is located at (and the filtering performed at) a client or a server.

By using the profile data elements discussed above and assigning privacy characteristics to each attribute-value pair, the user is able to make an informed tradeoff between the privacy of the profile data and the bandwidth and local storage requirements. For example, if the user has a strong privacy interest, then only a few of the attribute-value pairs may be assigned a "Public" privacy characteristic. In this example, less profile data is exposed to untrusted servers, so additional data is received and processed by the client. In another situation, if the user desires a reduction in bandwidth and local storage requirements, many of the attribute-value pairs may be assigned a "Public" privacy characteristic. In this situation, more profile data is exposed to untrusted servers, but less data is received and stored by the client.

The privacy characteristics associated with a particular attribute-value pair can be determined by the user or the data provider. A default privacy characteristic may be provided for some or all of the attribute-value pairs. For example, a default privacy characteristic of "Private" may be associated with all attribute-value pairs to avoid exposing any private information about the user unless the user specifically changes the default setting.

Embodiments of the invention allow users to further limit the distribution of attribute-value pairs to particular types of servers. For example, a user of a particular brand of computer may only want the "Model Number" attribute to be provided to servers associated with the manufacturer of the computer. Thus, the

-15-

“Model Number” may have a privacy characteristic of “Public”, but the attribute-value pair is only distributed to servers (or server filter criteria) associated with the particular manufacturer of the computer. The distribution of any attribute-value pair can be limited, regardless of the privacy characteristic. Additionally, a user may deactivate a particular attribute-value pair such that the attribute-value pair is not distributed to any server or client. The attribute-value pair remains deactivated until reactivated by the user. This deactivation provides a temporary way for a user to prevent filtering based on a particular attribute-value pair without permanently deleting the information from the profile data set.

Figures 8A and 8B illustrate exemplary server filter criteria and client filter criteria, respectively, generated from the profile data elements shown in Figure 7. The server filter criteria shown in Figure 8A contains two attribute-value pairs corresponding to the two "Public" entries shown in Figure 7. The server filter criteria shown in Figure 8A does not include the privacy characteristics. The privacy characteristics are used to determine which servers or clients will receive a particular attribute-value pair. However, the privacy characteristics are not transmitted along with the filter criteria.

Using the exemplary filter criteria shown in Figure 8A, a server is able to filter incoming data. For example, if the server receives incoming data (such as an advertisement or news article) targeted to male computer users over the age of 40, the filtering based on the server filter criteria will allow the data to pass to the next data filtering stage because the server filter criteria for John Doe identifies that John Doe is male. Although the next data filtering stage will reject the data because John Doe is not over 40, the server filter criteria does not include John Doe's age and therefore filtering using the server filter criteria cannot be based on that attribute. Using the example filter criteria shown in Figure 8A, the filtering based on the server filter criteria is only capable of filtering incoming data based on the user's name and gender. If the user changes the privacy characteristic associated with attribute “Age” to “Public,” then the server's filter criteria will include the attribute-value pair “Age – 38”. In this situation, the filtering based on the server filter criteria will filter out the incoming data based on John Doe's age.

-16-

Figure 8B contains six attribute-value pairs corresponding to the "Semi-Private" and "Private" entries shown in Figure 7. In this example, two filtering stages are used, but three levels of privacy characteristics are provided. Therefore, two of the privacy characteristic levels are combined into a single filtering stage. For this example, "Public" entries are provided in the server filter criteria and "Semi-Private" and "Private" entries are provided in the client filter criteria. In an alternative embodiment, the "Public" and "Semi-Private" entries are provided in the server filter criteria and the "Private" entries are provided in the client filter criteria. Although Figure 8B illustrates the client filter criteria separately from the profile data elements shown in Figure 7, embodiments of the invention may read the client filter criteria directly from the profile data elements instead of generating a separate instance of the client filter criteria.

Figures 8A and 8B illustrate server filter criteria and client filter criteria having distinct attributes; i.e., no shared attributes. Thus, the server filter criteria and the client filter criteria are completely different from one another. However, in other embodiments of the invention, one or more of the attributes may be contained in two or more filter criteria. For example, the attribute "Age" may be contained in both the server filter criteria and the client filter criteria such that both the server filtering stage and the client filtering stage perform data filtering using the "Age" attribute. However, the server filter criteria and the client filter criteria do not generally share all attributes. Any two filter criteria are "different" if at least one data element is different between the two criteria (e.g., a different attribute or a different attribute value).

Figure 9 illustrates another embodiment of a multi-stage data filtering system. The embodiment of Figure 9 represents a unified three-stage data filtering system (untrusted server, trusted server, and client). As mentioned above, the teachings of the present invention may be applied to data filtering systems having any number of data filtering stages. The components contained within the servers and the client in Figure 9 are similar to those discussed above with reference to Figure 4. Untrusted server 140 receives incoming data from a data source (not shown) and forwards identifier information regarding the data to client 144, via

-17-

trusted server 142. Client 144 filters the incoming data using the identifier information and an untrusted server filter criteria. An "untrusted server" request list for the data which satisfies the untrusted server filter criteria is then forwarded to untrusted server 140 via trusted server 142.

Upon receipt of the untrusted server request list, untrusted server 140 forwards the requested filtered data, if any, to trusted server 142. Trusted server 142 forwards identifier information regarding the data it received from untrusted server 140 to client 144. Client 144 filters the received identifier information using a trusted server filter criteria. A "trusted server" request list for the data which satisfies the trusted server filter criteria is then forwarded to trusted server 142. The filtered data, if any, is then communicated from trusted server 142 to client 144. Client 144 filters the received data using a client filter criteria to generate a final set of filtered data. The filtering process may be terminated at any point if the output of a particular filter removes all data.

Alternatively, client 144 may perform the filtering for multiple stages at the same time. Analogous to the discussion above regarding Figure 4, given that client 144 receives identifier information regarding all data received by untrusted server 140, client 144 may apply the trusted server filter criteria and/or the client filter criteria to the untrusted server request list generated from applying the untrusted server filter criteria while the untrusted server request list is being returned to untrusted server 140. Additionally, client 144 may apply the client filter criteria to the trusted server request list generated from applying the trusted server filter criteria to the untrusted server request list while the trusted server request list is being returned to trusted server 142.

Furthermore, analogous to the discussion above regarding Figure 4, the untrusted server filter criteria may be incorporated into the trusted server filter criteria (or the client filter criteria), and the trusted server filter criteria may be incorporated into the client filter criteria, thereby reducing the number of sets of filter criteria maintained by client 144.

Figure 10 illustrates another embodiment of a multi-stage data filtering system. The embodiment of Figure 10 represents a unified three-stage data

-18-

filtering system (untrusted server, trusted server, and client). As mentioned above, the teachings of the present invention may be applied to data filtering systems having any number of data filtering stages. The components contained within the servers and the client in Figure 10 are similar to those discussed above with reference to Figure 9. Untrusted server 150 receives incoming data from a data source (not shown) and forwards identifier information regarding the data to trusted server 152. Trusted server 152 filters the incoming data using the identifier information and an untrusted server filter criteria. An untrusted server request list for the data which satisfies the untrusted server filter criteria is then returned to untrusted server 150.

Upon receipt of the request list, untrusted server 150 forwards the requested filtered data, if any, to trusted server 152. Trusted server 152 forwards identifier information regarding the data it received from untrusted server 150 to client 154. Client 154 filters the received identifier information using a trusted server filter criteria. A trusted server request list for the data which satisfies the trusted server filter criteria is then returned to trusted server 152. The filtered data, if any, is then communicated from trusted server 152 to client 154. Client 154 filters the received data using a client filter criteria to generate a final set of filtered data. The filtering process may be terminated at any point if the output of a particular filter removes all data.

Alternatively, multiple steps of the multiple-stage filtering in Figure 10 may be performed at the same time. Similar to the discussion above regarding Figure 4, identifier information corresponding to the information on the untrusted server request list may be transmitted to client 154 while the untrusted server request list is being returned to server 150. This allows client 154 to begin filtering, using the trusted server filter criteria and/or the client filter criteria, while the data from the untrusted server request list is being transferred from untrusted server 150 to trusted server 152.

Furthermore, analogous to the discussion above regarding Figure 4, the trusted server filter criteria may be incorporated into the client filter criteria, thereby reducing the number of sets of filter criteria maintained by client 154.

Figure 11 illustrates another embodiment of a multi-stage data filtering system. The embodiment of Figure 11 represents a unified three-stage data filtering system (untrusted server, trusted server, and client). As mentioned above, the teachings of the present invention may be applied to data filtering systems having any number of data filtering stages. The components contained within the servers and the client in Figure 11 are similar to those discussed above with reference to Figure 1. Untrusted server 160 receives incoming data from a data source (not shown) and filters the incoming data using an untrusted server filter criteria. The filtered data, if any, is then communicated from untrusted server 160 to trusted server 162. Trusted server 162 filters the received data using a trusted server filter criteria. The filtered data, if any, is then communicated from trusted server 162 to client 164. Client 164 filters the received data using a client filter criteria to generate a final set of filtered data. The filtering process may be terminated at any point if the output of a particular filter removes all data.

Figures 12A - 12C illustrate exemplary filter criteria for use in the three-stage data filtering system shown in Figure 11. Figures 12A - 12C use the exemplary profile data elements shown in Figure 7. Figure 12A illustrates an untrusted server filter criteria (i.e., the attribute-value pairs having a privacy characteristic "Public"). Figure 12B illustrates a trusted server filter criteria (i.e., the attribute-value pairs having a "Semi-Private" privacy characteristic). Figure 12C illustrates a client filter criteria (i.e., the attribute-value pairs having a privacy characteristic "Private").

Figure 13 illustrates another embodiment of a multi-stage data filtering system in which a client 186 receives data from multiple servers 170-184. A single profile data set is stored in client 186. Client 186 distributes various attribute-value pairs to the filtering criteria for multiple servers, located at client 186, based on the trustworthiness of the server and the privacy characteristics associated with each attribute-value pair. For example, for untrusted servers 170, 172, and 174, client 186 may include an untrusted server filter criteria containing only "Public" attribute-value pairs, and for trusted server 180 client 186 may include a trusted server filter criteria containing "Semi-Private" attribute-value

-20-

pairs. Additionally, for trusted server 184, client 186 may include a trusted server filter criteria containing "Public" and "Semi-Private" attribute-value pairs. For untrusted server 172, client 186 may include "Public" attribute-value pairs, while the "Semi-Private" and "Private" attribute pairs are filtered using the client filter criteria. Thus, client 186 may be filtering, using the client filter criteria, "Private" attribute-value pairs for some incoming data and filtering, using the client filter criteria, "Semi-Private" and "Private" attribute-value pairs for other incoming data.

It is not necessary that data filtering occur corresponding to every device through which the data passes. For example, for untrusted servers 176 and 178, client 186 may include "Public" attribute-value pairs, and include the remaining "Semi-Private" and "Private" attribute-value pairs in the client filter criteria. In this example, the filtered data from untrusted servers 176 and 178 passes through trusted server 182 without any data filtering operation based on trusted server filter criteria.

Figure 14 illustrates an embodiment of a computer system that can be used with the present invention (e.g., as a client or a server). The various components shown in Figure 14 are provided by way of example. Certain components of the computer in Figure 14 can be deleted from the data filtering system for a particular implementation of the invention, or additional components can be added (e.g., additional processors, buses, I/O devices, etc.). The computer shown in Figure 14 may be any type of computer including a general purpose computer.

Figure 14 illustrates a system bus 190 to which various components are coupled. A processor 192 performs the processing tasks required by the computer. Processor 192 may be any type of processing device capable of implementing the steps necessary to perform the data filtering operations discussed above. An input/output (I/O) device 194 is coupled to bus 190 and provides a mechanism for communicating with other devices coupled to the computer. A read-only memory (ROM) 196 and a random access memory (RAM) 198 are coupled to bus 190 and provide a storage mechanism for various data and information used by the computer. Although ROM 196 and RAM 198 are shown coupled to bus 190, in alternate embodiments, ROM 196 and RAM

-21-

198 are coupled directly to processor 192 or coupled to a dedicated memory bus (not shown).

A video display 200 is coupled to bus 190 and displays various information and data to the user of the computer. A disk drive 202 is coupled to bus 190 and provides for the long-term mass storage of information. Disk drive 202 may be used to store various profile data sets and other data generated by and used by the data filtering system. A keyboard 204 and pointing device 208 are also coupled to bus 190 and provide mechanisms for entering information and commands to the computer. A printer 206 is coupled to bus 190 and is capable of creating a hard-copy of information generated by or used by the computer.

Figure 15 illustrates an embodiment of a computer-readable medium 220 containing various sets of instructions, code sequences, configuration information, and other data used by a computer or other processing device. The embodiment illustrated in Figure 15 is suitable for use with the data filtering system described above. The various information stored on medium 220 is used to perform various data filtering and data processing operations. Computer-readable medium 220 is also referred to as a processor-readable or machine-readable medium. Computer-readable medium 220 can be any type of magnetic, optical, or electrical storage medium including a diskette, magnetic tape, CD-ROM, DVD, memory device, or other storage medium.

Computer-readable medium 220 includes interface code 222 that controls the flow of information between various devices or components in a data filtering system. Interface code 222 may control the transfer of information within a device (e.g., between the processor and a memory device), or between an input/output port and a storage device. Additionally, interface code 222 may control the transfer of information from one device to another (e.g., the transfer of filtered data or profile data between a client and a server). Data filtering code 234 filters received data based on a particular filter criteria, as discussed above.

Computer-readable medium 220 also includes a profile data set 224 used to filter data and generate filter criteria. Profile data set 224 may include user-specific information, information related to user role(s), and/or information related to user

-22-

class(es). Filter criteria 236 is used by the data filtering procedures described above, and may include multiple sets of filter criteria for multiple stages (e.g., untrusted server filter criteria, trusted server filter criteria, and client filter criteria). Received data 226 represents data and/or identifier information that has been received by a particular device for filtering. Received data 226 may be filtered data from another device, may be unfiltered incoming data distributed by a third-party data source, or may be identifier information from another device. Filtered data 238 represents the output of the data filtering process as applied to received data 226. If the filtering process filters out (i.e., removes) all received data 226, then filtered data 238 may be a null set.

Profile data generation code 228 typically resides on a client, and is used to generate profile data set 224. Profile data generation code 228 may be executed by a user of the client to generate or modify the various profile data attributes, values, and privacy characteristics contained in profile data set 224. Computer-readable medium 220 also includes code 240 for determining a level of trust associated with a particular device (such as a server). Typically, this code 240 is executed by a user of the client and may assign a default level of trust to a particular device if a level of trust is not otherwise assigned. For example, a default level of trust may be "untrusted," such that the device only receives profile data having a privacy characteristic of "Public."

Filtered data processing code 230 processes filtered data 238. For example, data processing code 230 may display filtered data 238 to a user, notify a user of the received data, or communicate filtered data 238 to the next device (e.g., transmit filtered data 238 from a server to a client). Filter criteria generation code 242 generates filter criteria based on information contained in profile data set 224 and the level of trust for a particular device as determined by code 240. Typically, filter generation code 242 is executed by a client, which generates a filter criteria for a particular device. The filter criteria contains the attributes and values from profile data set 224 that correspond to the level of trust associated with the particular device. For example, an untrusted server may only receive attributes and values having a privacy characteristic of "Public." Therefore, the filter criteria

-23-

for an untrusted server will not contain attributes and values having a privacy characteristic of "Semi-Private" or "Private."

Computer-readable medium 220 also includes information 232 regarding user role(s) and information 244 regarding user class(es). As discussed above, information relating to user roles and user classes identify characteristics or attributes associated with roles or classes, rather than an individual person. As shown in Figure 15, information 232 regarding user role(s) and information 244 regarding user class(es) may be stored separately from profile data set 224. In alternate embodiments, information regarding user role(s) and class(es) may be stored within profile data set 224.

Figure 15 illustrates an exemplary computer-readable medium 220 containing various sets of instructions, code sequences, and other information that can be used by a data filtering system. However, in particular data filtering devices, one or more of the items illustrated in Figure 15 may not be required. For example, in a computer-readable medium for use with an untrusted server that relies on a client for its filter criteria 236, the computer-readable medium need not contain profile data set 224, profile data generation code 228, code 240 for determining level of trust, filter criteria generation code 242, or information 232 and 244 regarding user role(s) and user class(es). In this example, the client maintains the profile data set, generates the filter criteria for the untrusted server, and communicates the filter criteria to the untrusted server. To maintain the privacy of the profile data set, the profile data set is typically stored only on the client.

In some of the embodiments discussed above, reference is made to a server forwarding identifier information to another server or client. In alternate embodiments, additional parts of the data may also be forwarded. For example, in situations where the size (e.g., number of bytes) of the data is not much larger than the size of the identifier information, then all of the data may be forwarded rather than just the identifier information.

Additionally, specific examples are given above using two stages (server and client filtering) as well as three stages (untrusted server, trusted server, and

-24-

client filtering), where more trustworthy devices perform filtering for themselves and/or one or more additional less trustworthy devices. However, in alternate embodiments of the present invention, any number of stages of filtering can be carried out by a device for any number of less trustworthy devices.

Thus, a multi-stage data filtering system has been described that does not compromise a user's privacy. The system provides a filtering system that advantageously distributes multiple profile data elements to two or more data filtering stages, even though multiple data filtering stages may be performed by a single device or system.

From the above description and drawings, it will be understood by those skilled in the art that the particular embodiments shown and described are for purposes of illustration only and are not intended to limit the scope of the invention. Those skilled in the art will recognize that the invention may be embodied in other specific forms without departing from its spirit or essential characteristics. References to details of particular embodiments are not intended to limit the scope of the claims.

-25-

CLAIMS

What is claimed is:

1. A method of filtering data in a first device, the method comprising:
receiving, from a second device, identifier information corresponding to a first set of data;
filtering the first set of data based on a first filter criteria and the identifier information, wherein the filtering of the first set of data identifies a first set of filtered data;
receiving, subsequent to the filtering of the first set of data, the first set of filtered data from the second device; and
filtering the first set of filtered data based on a second filter criteria, wherein the filtering of the first set of filtered data generates a second set of filtered data, and wherein the second filter criteria is different from the first filter criteria.
2. The method of claim 1 wherein the first filter criteria and the second filter criteria are included in a profile data set.
3. The method of claim 2 wherein the first filter criteria contains public profile data.
4. The method of claim 2 wherein the second filter criteria contains private profile data.
5. The method of claim 2 wherein the profile data set is associated with a particular data recipient.
6. The method of claim 1 wherein the second device is an untrusted filtering device and the first device is a trusted filtering device.

-26-

7. The method of claim 1 further including displaying the second set of filtered data.

8. The method of claim 1 wherein the second device is a server device and the first device is a client device.

9. A method of filtering data in a first device, the method comprising:
filtering data based on both a first set of identifiers corresponding to the data and a first filter criteria, to generate a first set of filtered data identifiers; and
filtering at least a portion of the first set of filtered data identifiers based on a second filter criteria, to generate a second set of filtered data identifiers.

10. The method of claim 9 wherein the first filter criteria and the second filter criteria contain different filter characteristics.

11. The method of claim 9 wherein the first filter criteria contains public profile data and the second filter criteria contains private profile data.

12. The method of claim 11 further comprising:
providing the first set of filtered data identifiers to a second device;
providing the second set of filtered data identifiers to a third device; and
receiving filtered data corresponding to the second set of filtered data identifiers from the third device.

13. A machine-readable medium having stored thereon a plurality of instructions, designed to be executed by a processor, for implementing a function to filter data in a first device, the plurality of instructions including instructions to:
receive, from a second device, identifier information corresponding to a first set of data;

-27-

filter the first set of data based on a first filter criteria and the identifier information, wherein the filtering of the first set of data identifies a first set of filtered data;

receive, subsequent to the filtering of the first data, the first set of filtered data from the second device; and

filter the first set of filtered data based on a second filter criteria, wherein the filtering of the first set of filtered data generates a second set of filtered data, and wherein the second filter criteria is different from the first filter criteria.

14. The machine-readable medium of claim 13 wherein the second device is a server device and the first device is a client device.

15. A machine-readable medium having stored thereon a plurality of instructions, designed to be executed by a processor, for implementing a function to filter data in a first device, the plurality of instructions including instructions to:

filter data based on both a first set of identifiers corresponding to the data and a first filter criteria, to generate a first set of filtered data identifiers; and

filter at least a portion of the first set of filtered data identifiers based on a second filter criteria, to generate a second set of filtered data identifiers.

16. The machine-readable medium of claim 15, wherein the plurality of instructions further include instructions to:

provide the first set of filtered data identifiers to a second device;

provide the second set of filtered data identifiers to a third device; and

receive filtered data corresponding to the second set of filtered data identifiers from the third device.

17. An apparatus for filtering data, the apparatus comprising:
an interface to receive, from a device, identifier information corresponding to a first set of data, and to subsequently receive a first set of filtered data from the device;
and

-28-

a filter, coupled to the interface, to filter the first set of data based on a first filter criteria and the identifier information, wherein the filtering of the first set of data identifies the first set of filtered data, and to filter the first set of filtered data based on a second filter criteria, wherein the filtering of the first set of filtered data generates a second set of filtered data, and wherein the second filter criteria is different from the first filter criteria.

18. The apparatus of claim 17 wherein the device is a server device and the apparatus is a client device.

19. An apparatus for filtering data, the apparatus comprising:
an interface to communicate with other devices; and
a filter, coupled to the interface, to filter data based on both a first set of identifiers corresponding to the first data and a first filter criteria in order to generate a first set of filtered data identifiers, and to filter at least a portion of the first set of filtered data identifiers based on a second filter criteria in order to generate a second set of filtered data identifiers.

20. The apparatus of claim 19 wherein the interface is further to provide the first set of filtered data identifiers to a second device, provide the second set of filtered data identifiers to a third device, and receive filtered data corresponding to the second set of filtered data identifiers from the third device.

1 / 14

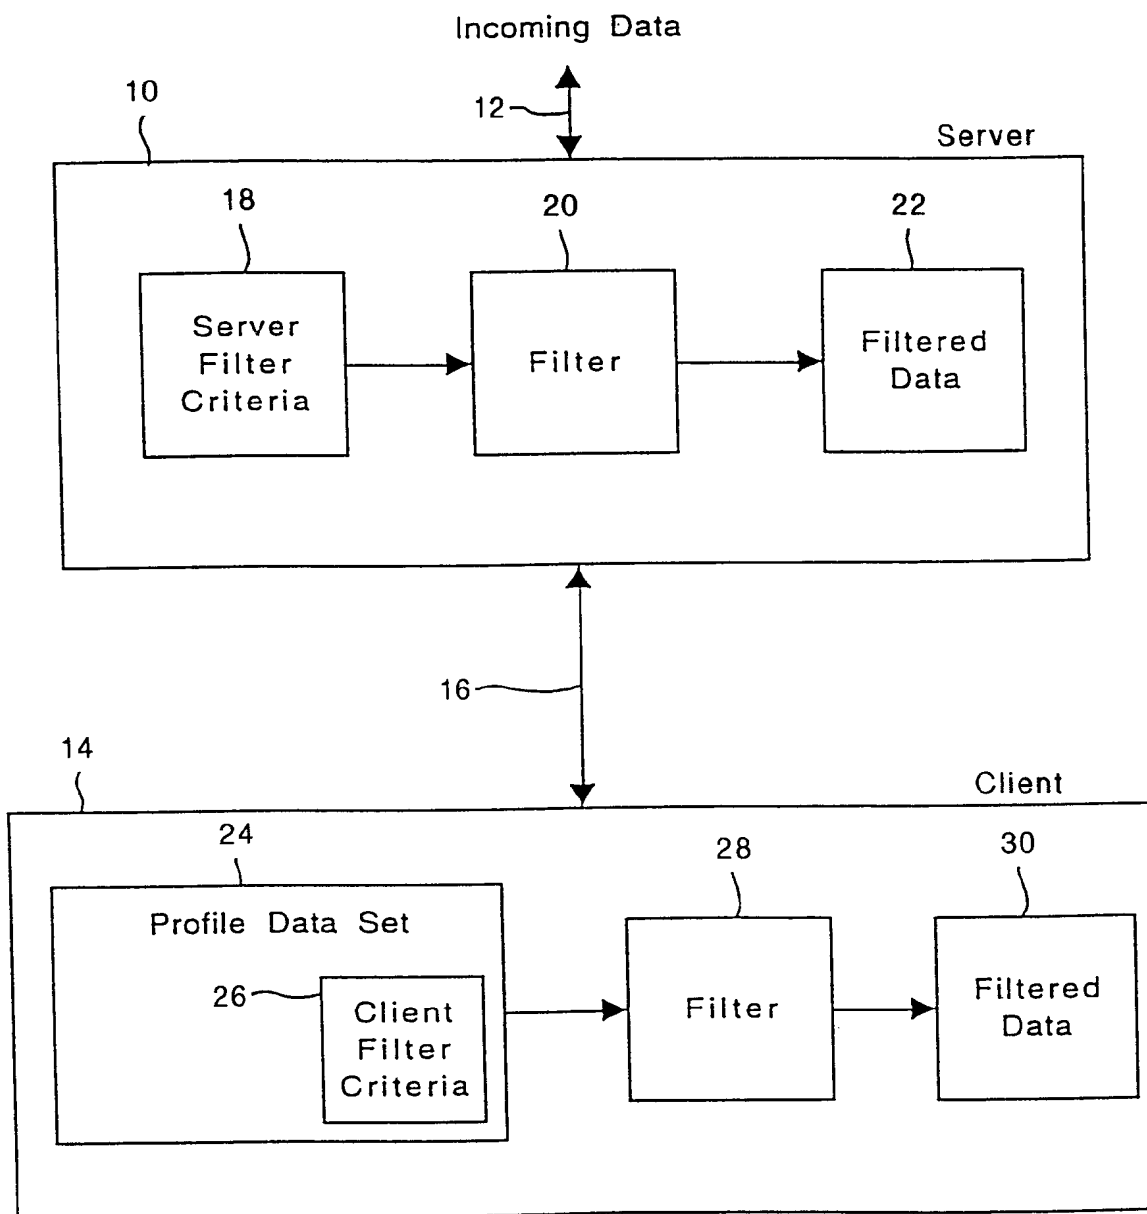


FIG. 1

2 / 14

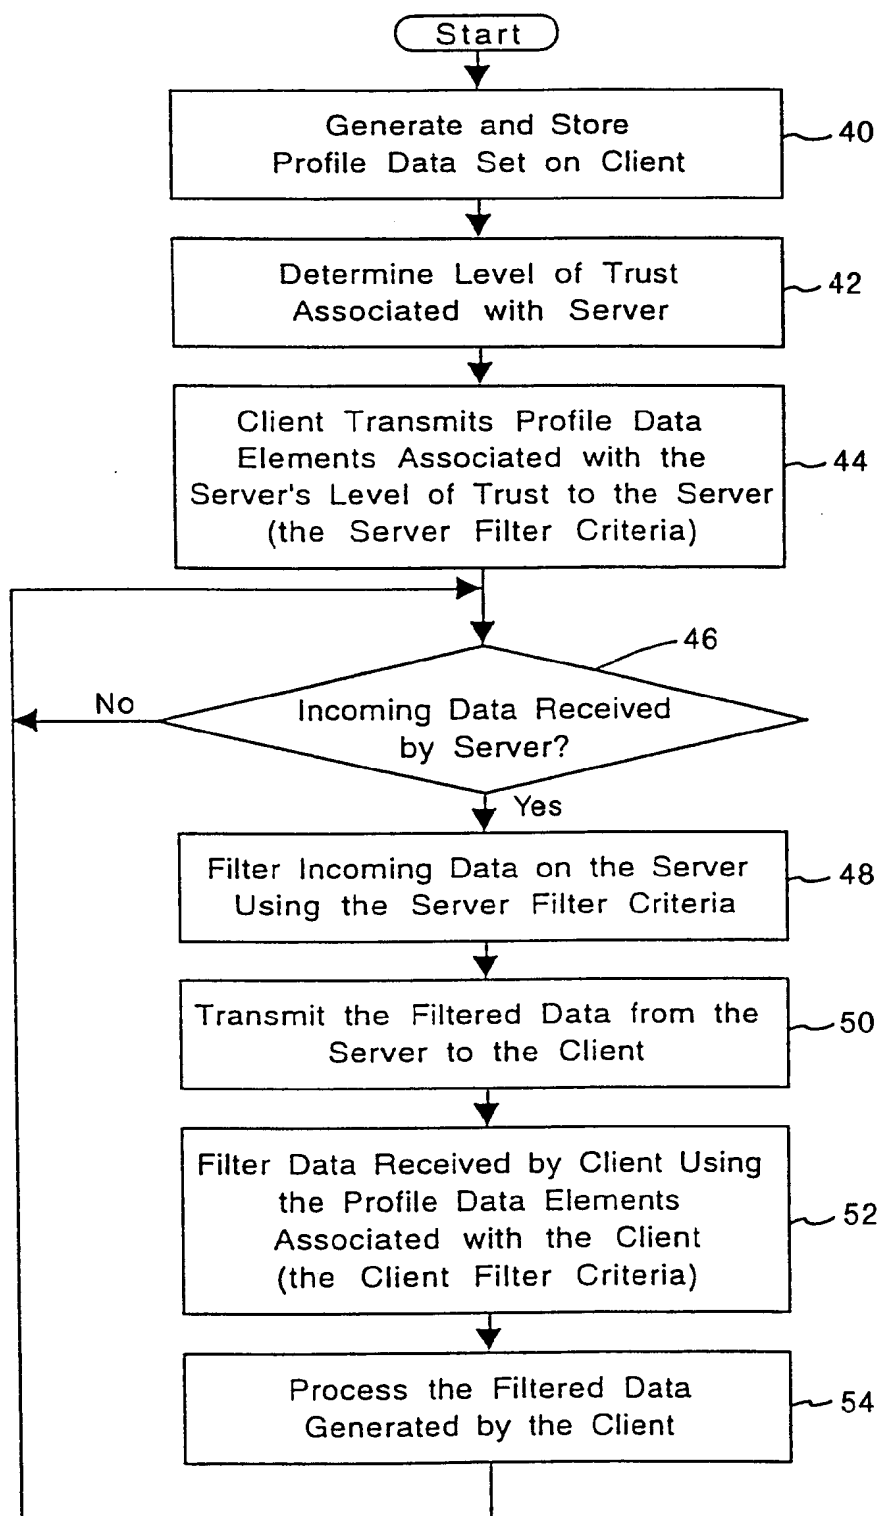


FIG. 2

3 / 14

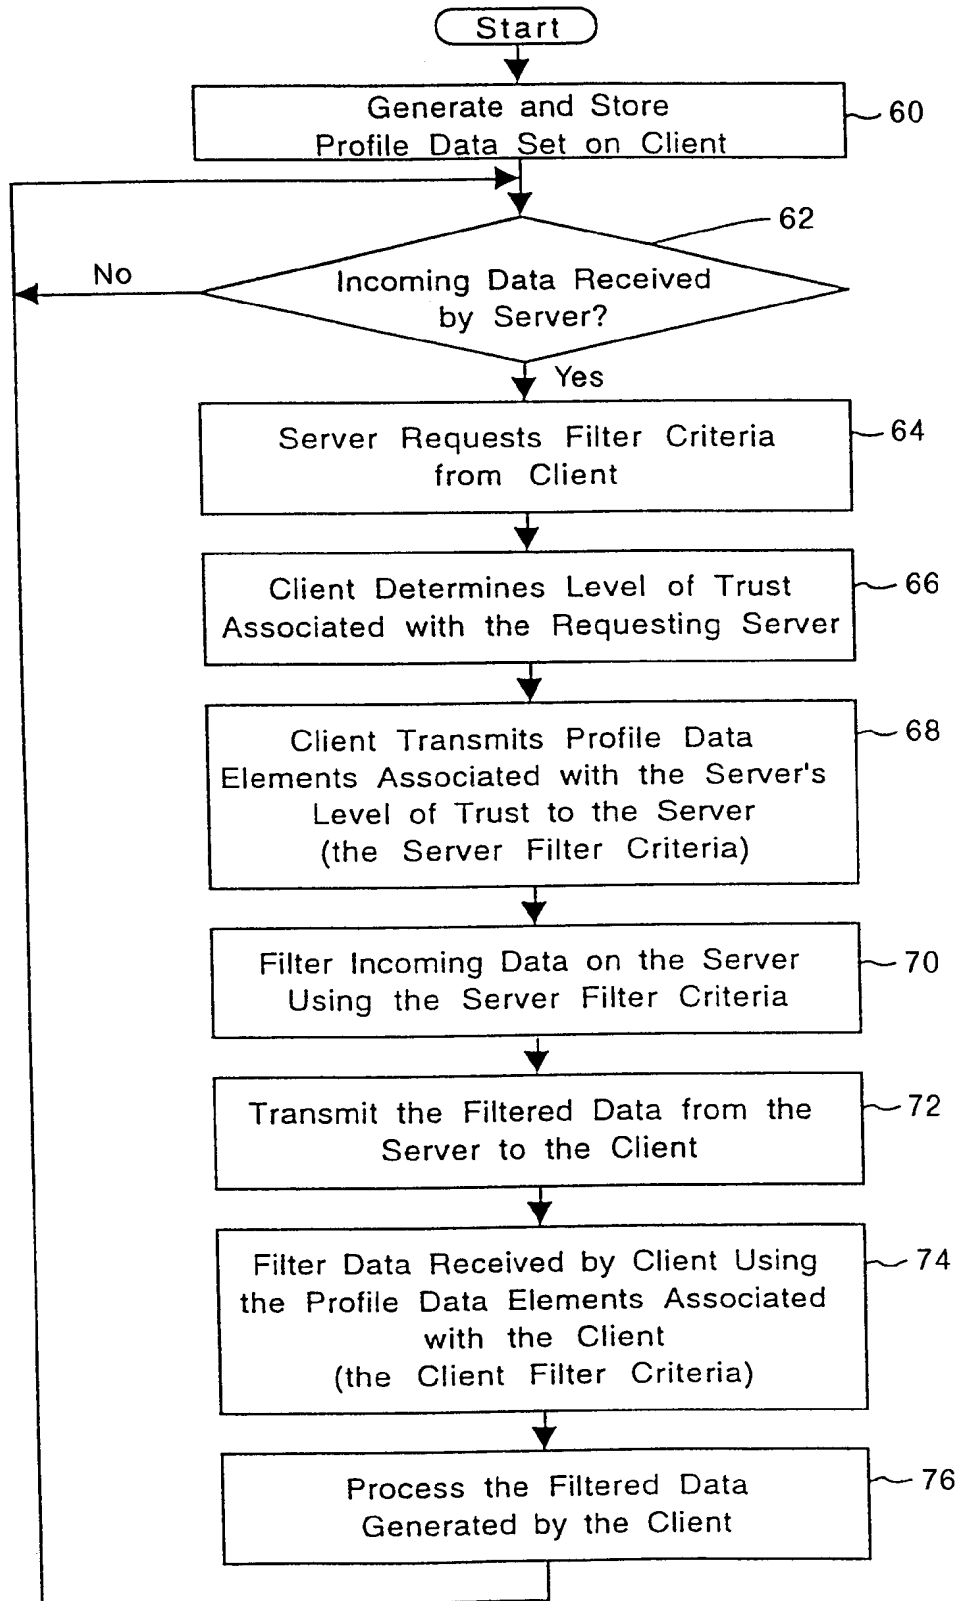


FIG. 3
SUBSTITUTE SHEET (RULE 26)

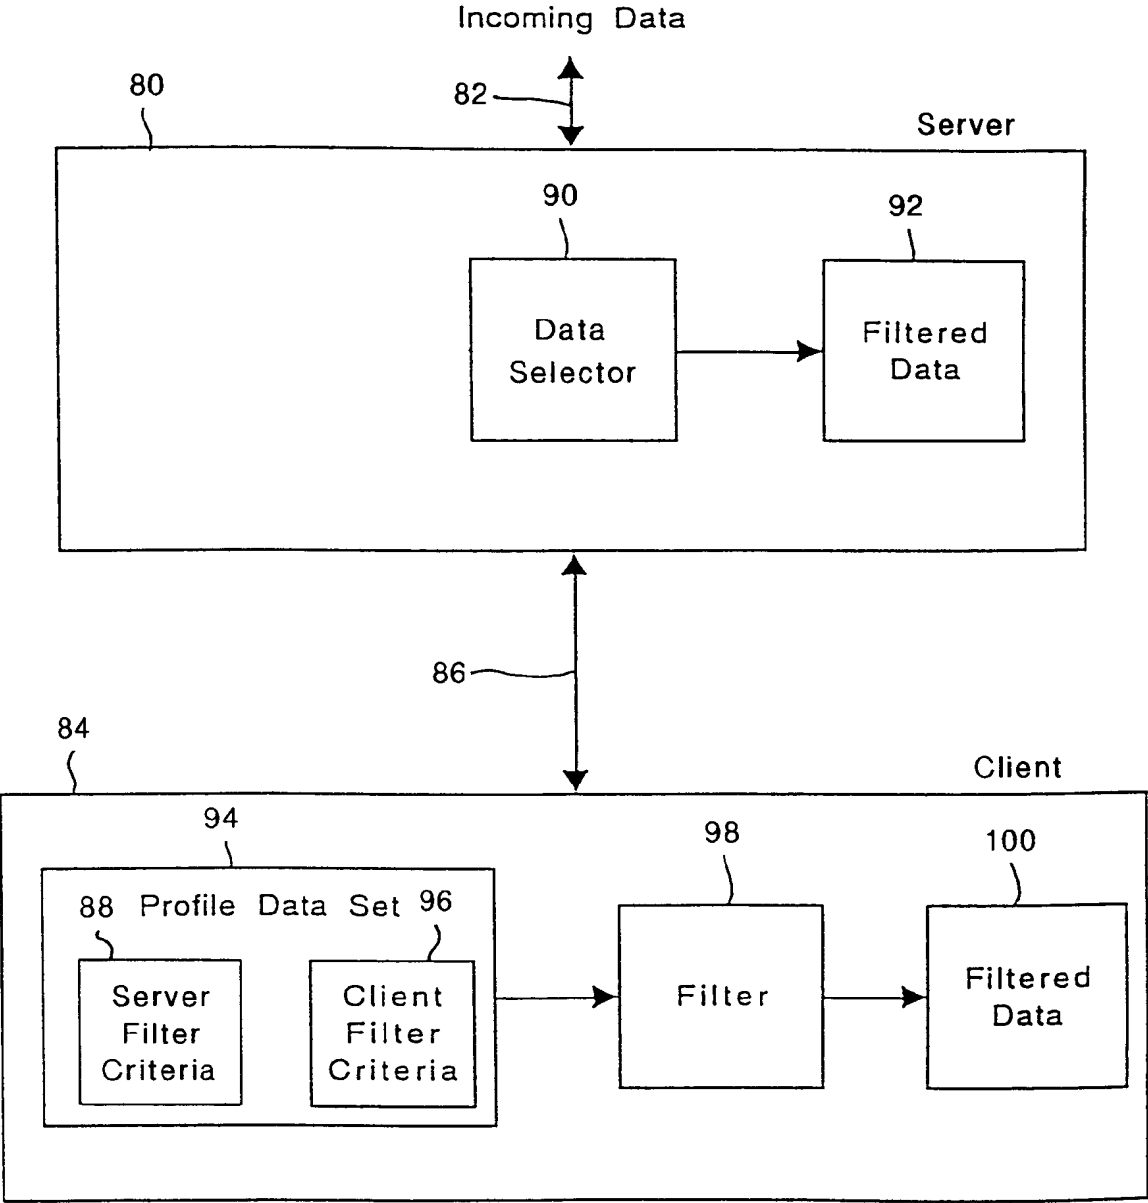


FIG. 4

5 / 14

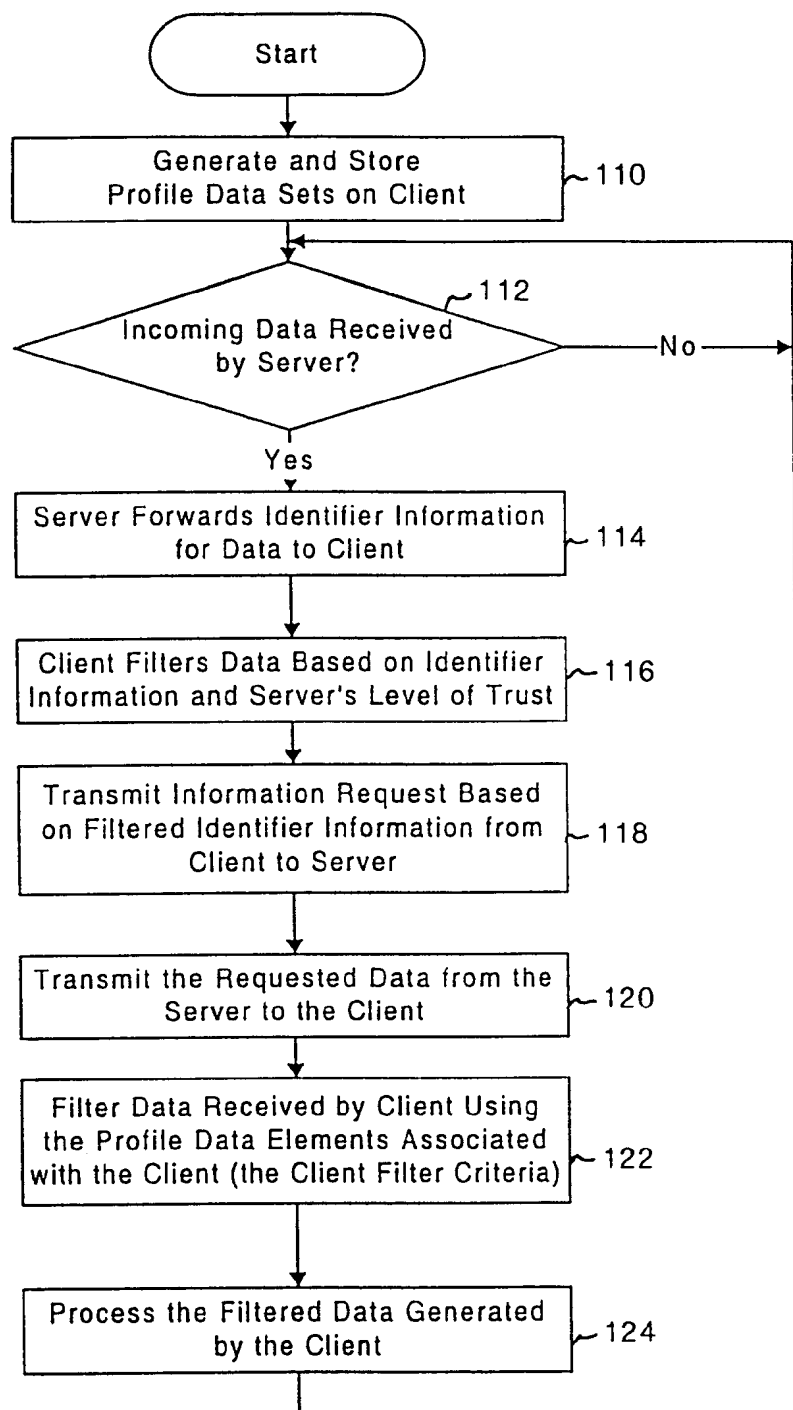
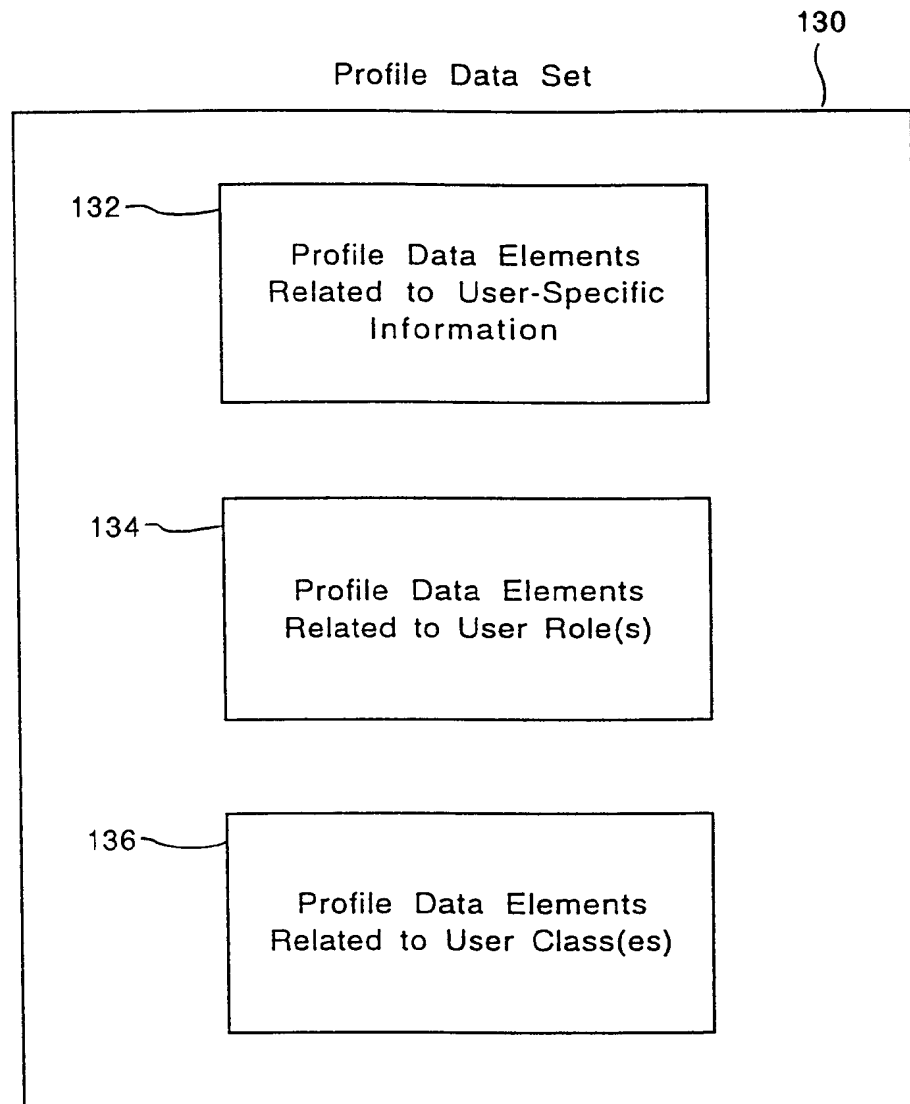


FIG. 5

6 / 14

**FIG. 6**

7 / 14

Attribute	Value	Privacy Characteristics
Name	John Doe	Public
Employer	Acme Corp.	Semi-Private
Occupation	Engineer	Semi-Private
Age	38	Semi-Private
Gender	Male	Public
Height	6-0	Semi-Private
Social Security No.	123-45-6789	Private
Marital Status	Married	Private

FIG. 7

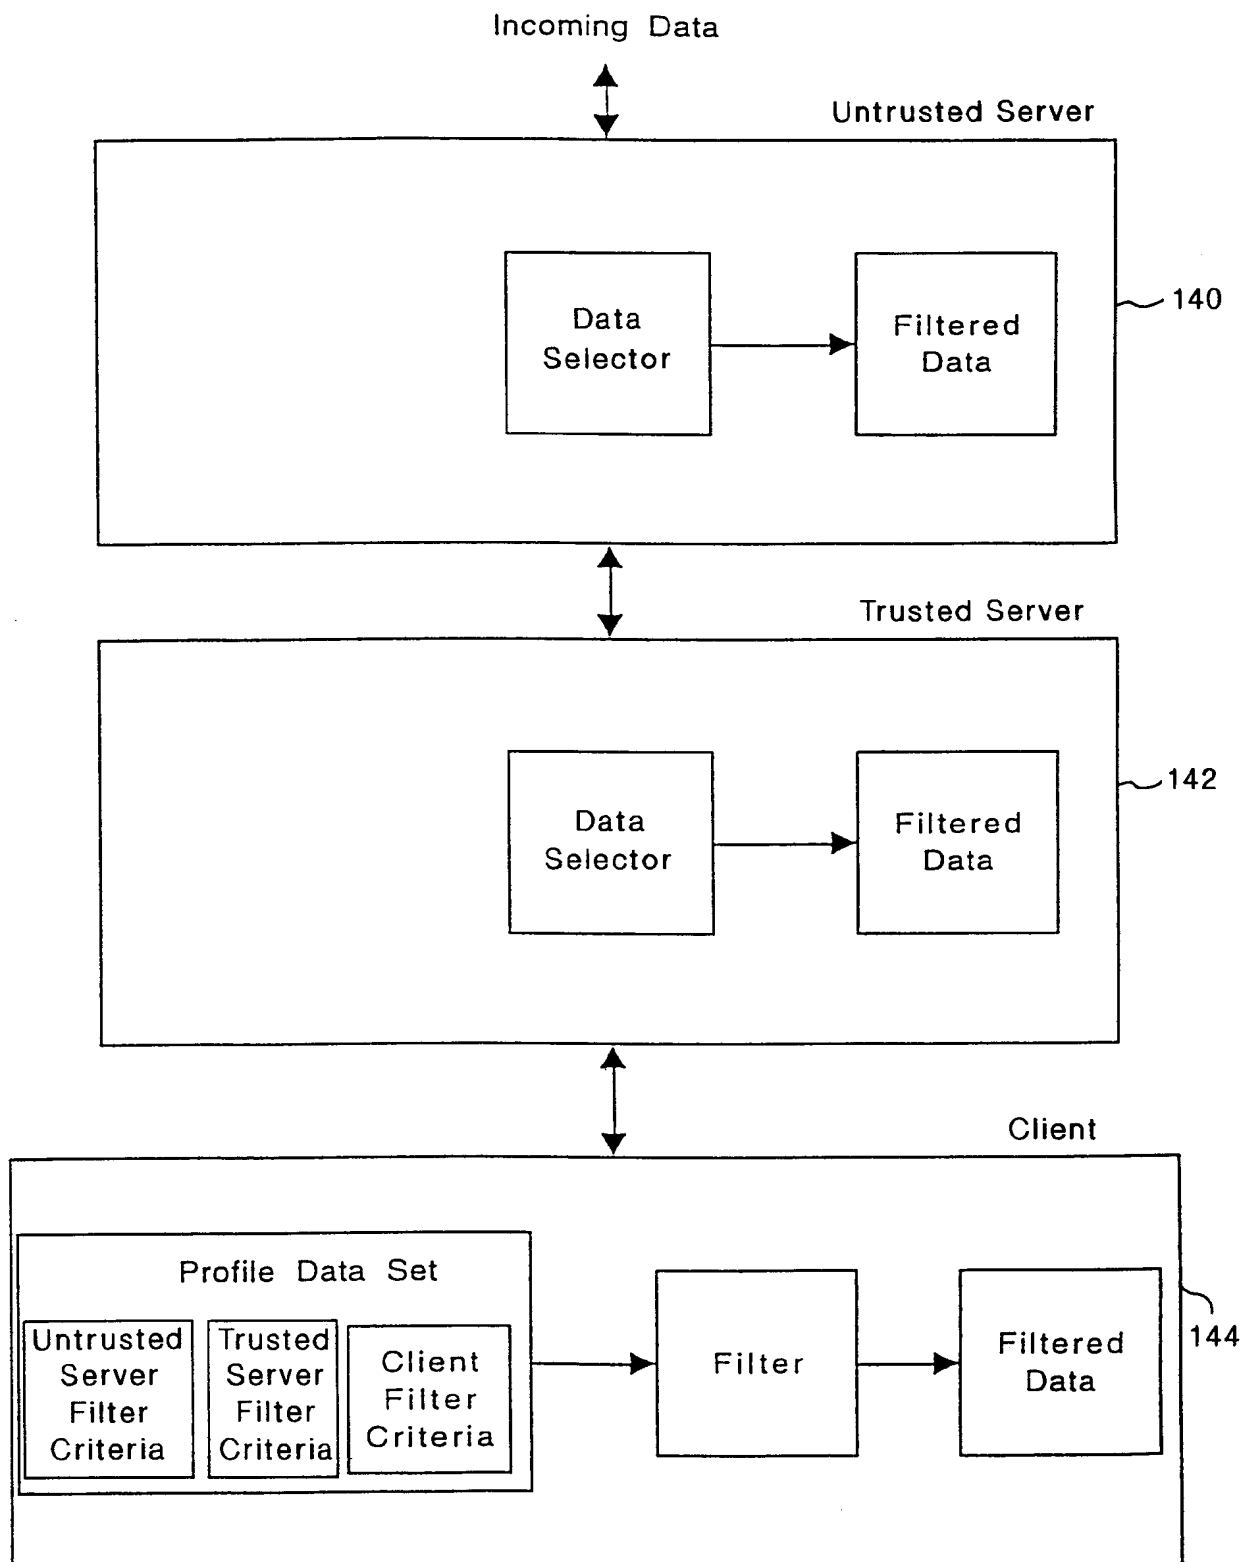
Attribute	Value
Name	John Doe
Gender	Male

FIG. 8A

Attribute	Value
Employer	Acme Corp.
Occupation	Engineer
Age	38
Height	6-0
Social Security No.	123-45-6789
Marital Status	Married

FIG. 8B

8 / 14

**FIG. 9**

9 / 14

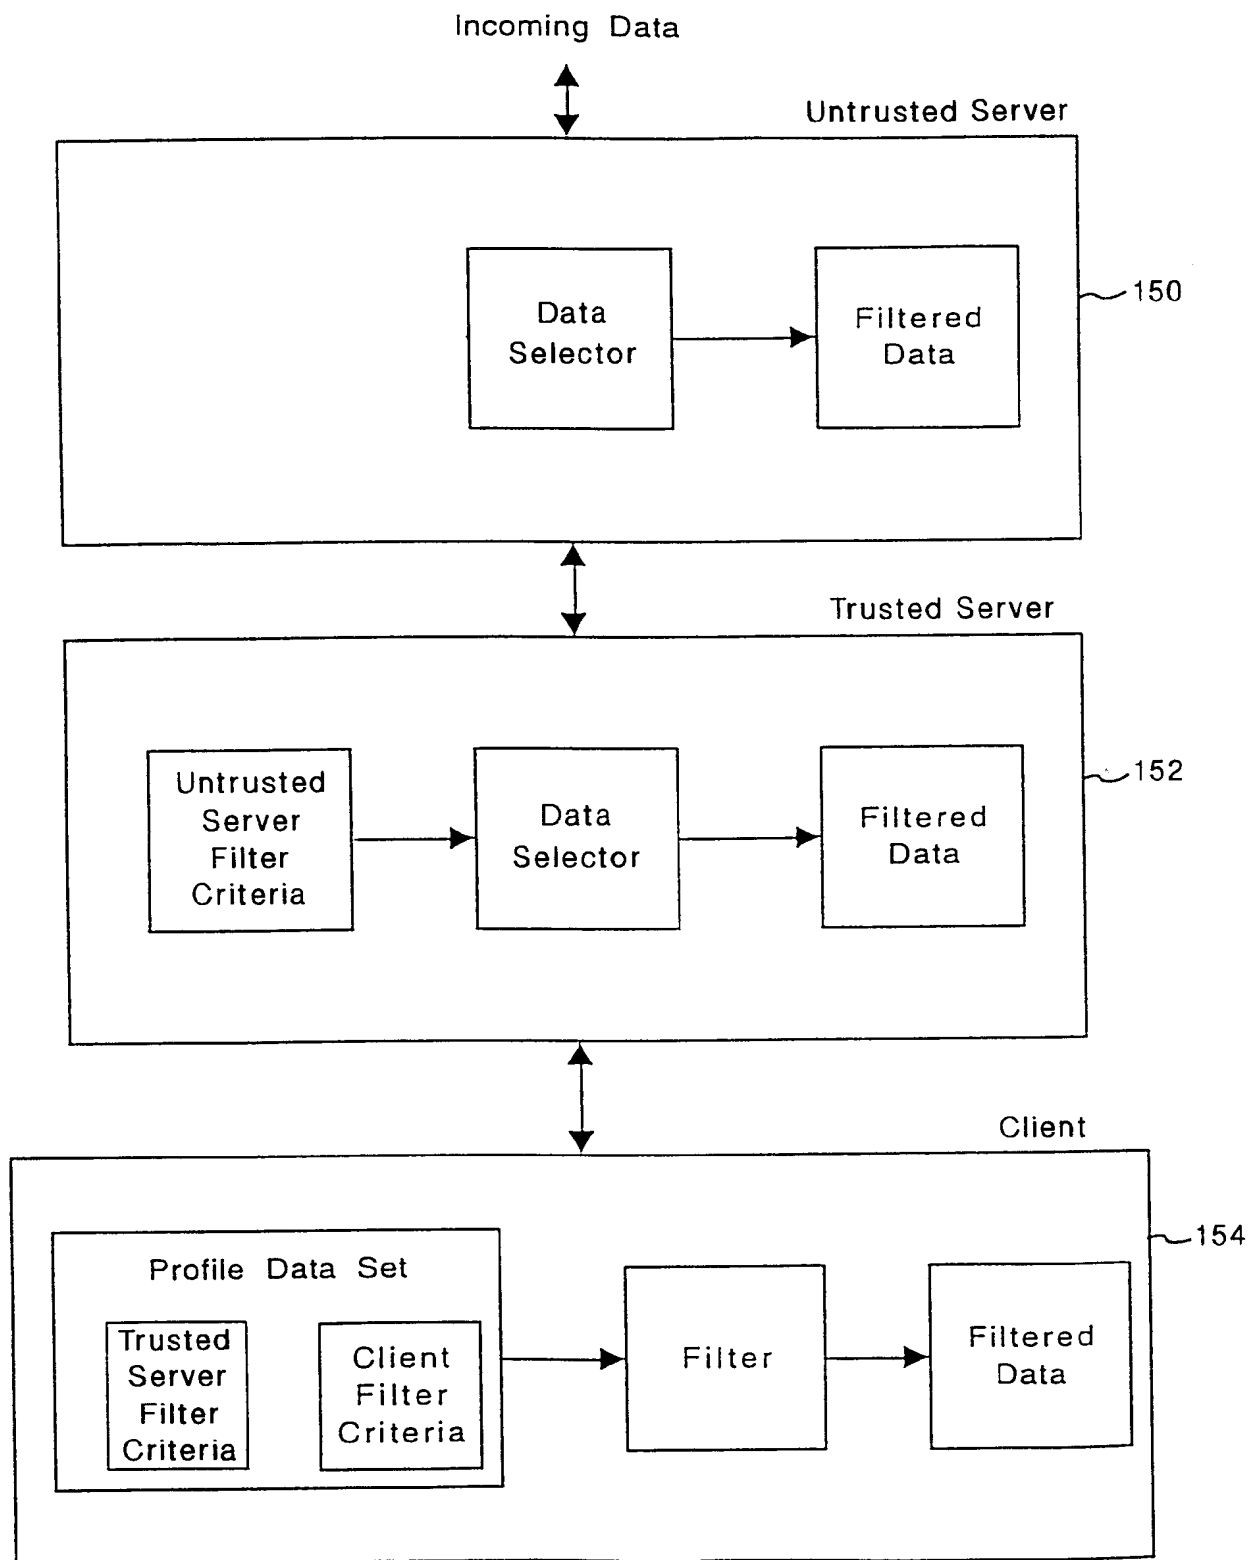


FIG. 10

10 / 14

Incoming Data




FIG. 11

Attribute	Value
Name	John Doe
Gender	Male

FIG. 12A

Attribute	Value
Employer	Acme Corp.
Occupation	Engineer
Age	38
Height	6-0

FIG. 12B

Attribute	Value
Social Security No.	123-45-6789
Marital Status	Married

FIG. 12C

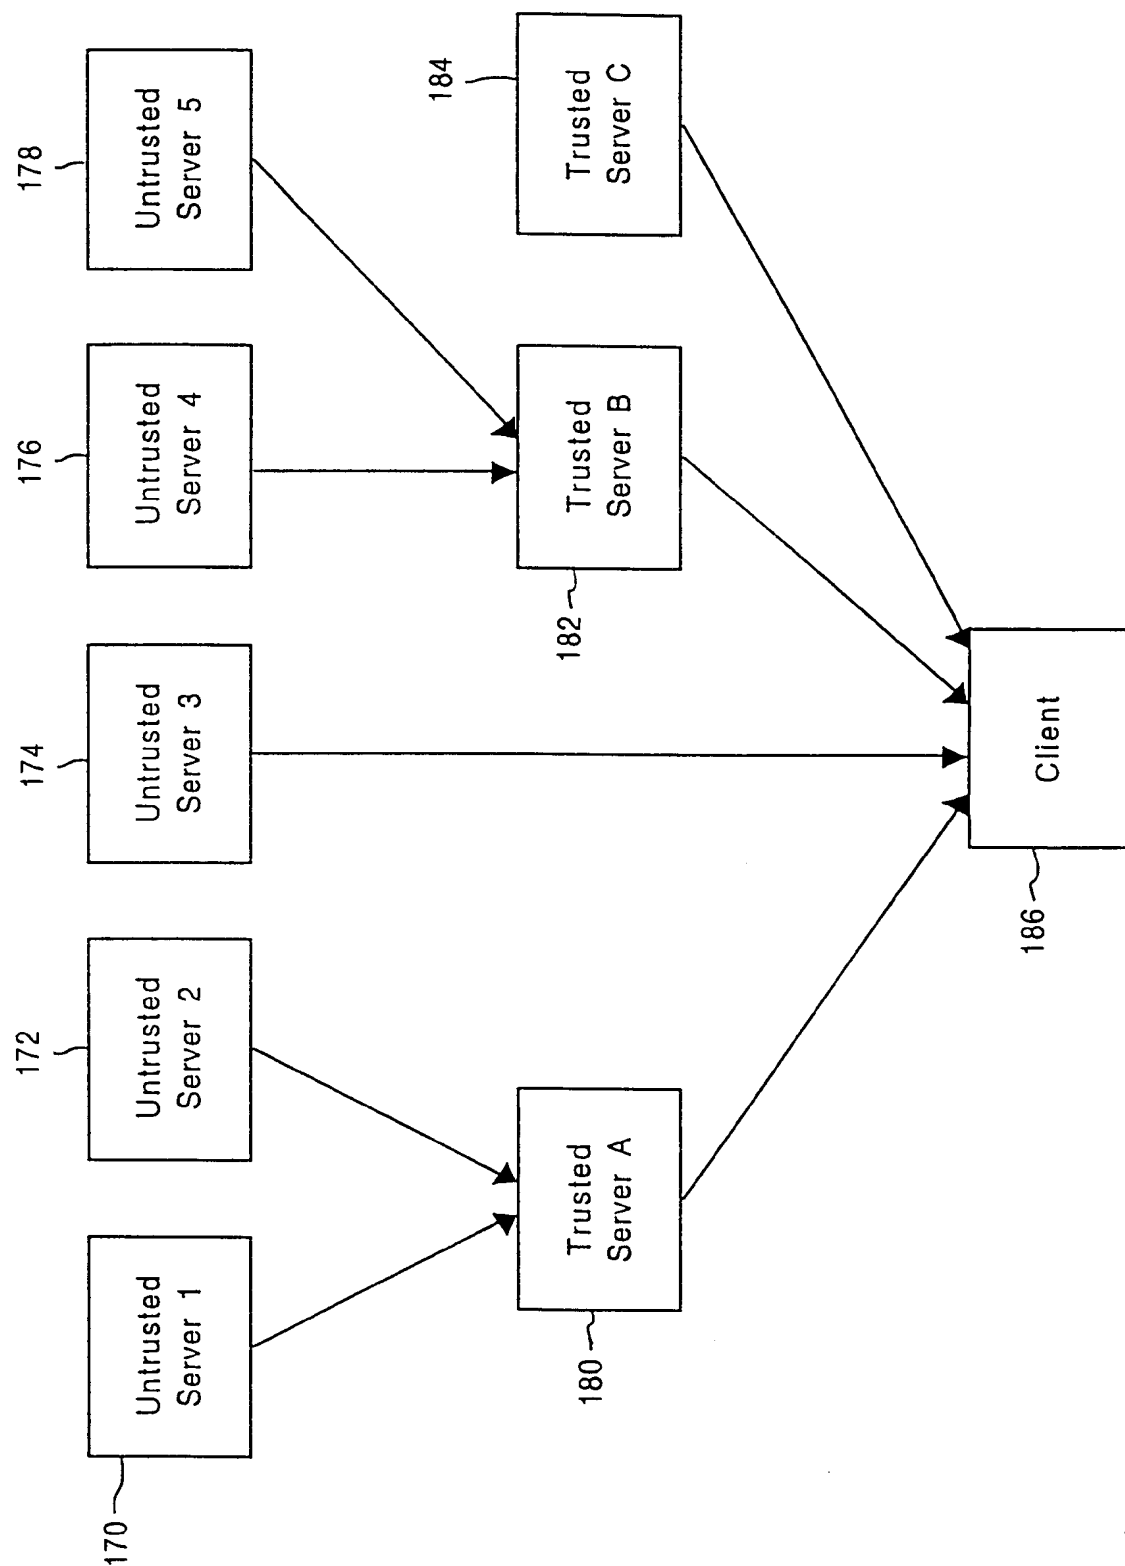


FIG. 13

13 / 14

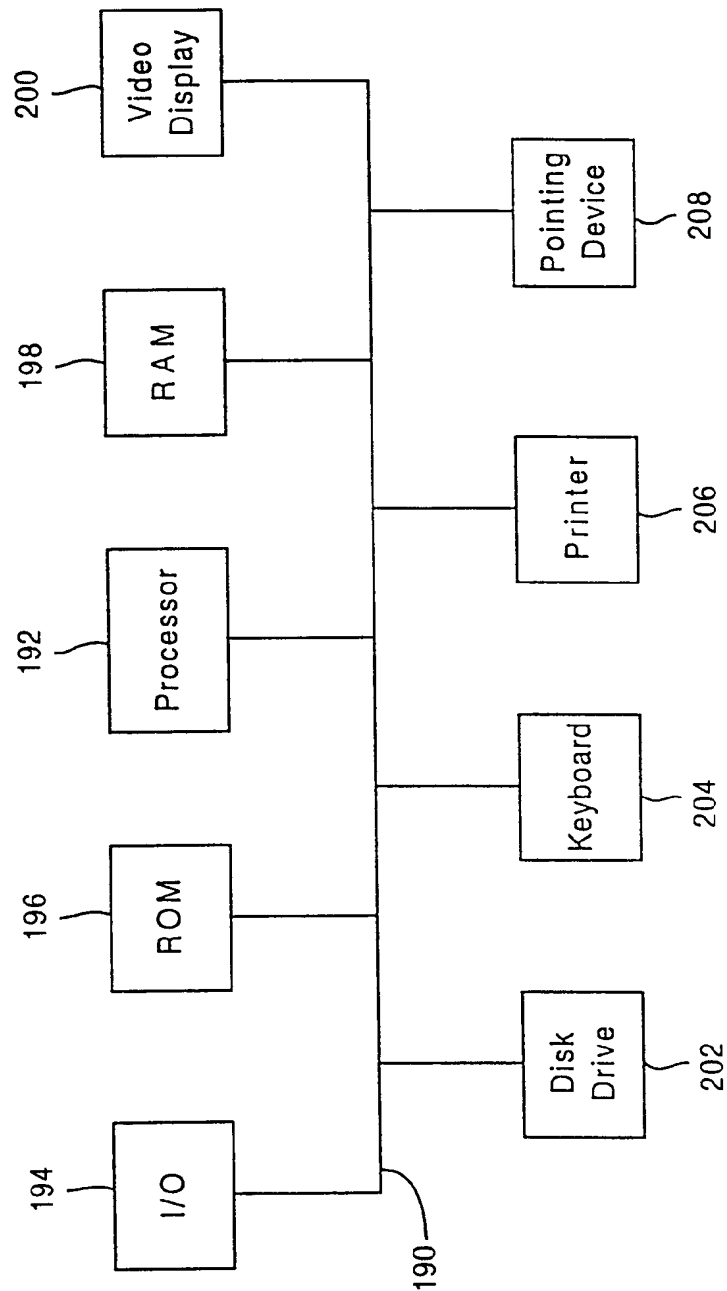


FIG. 14

14 / 14

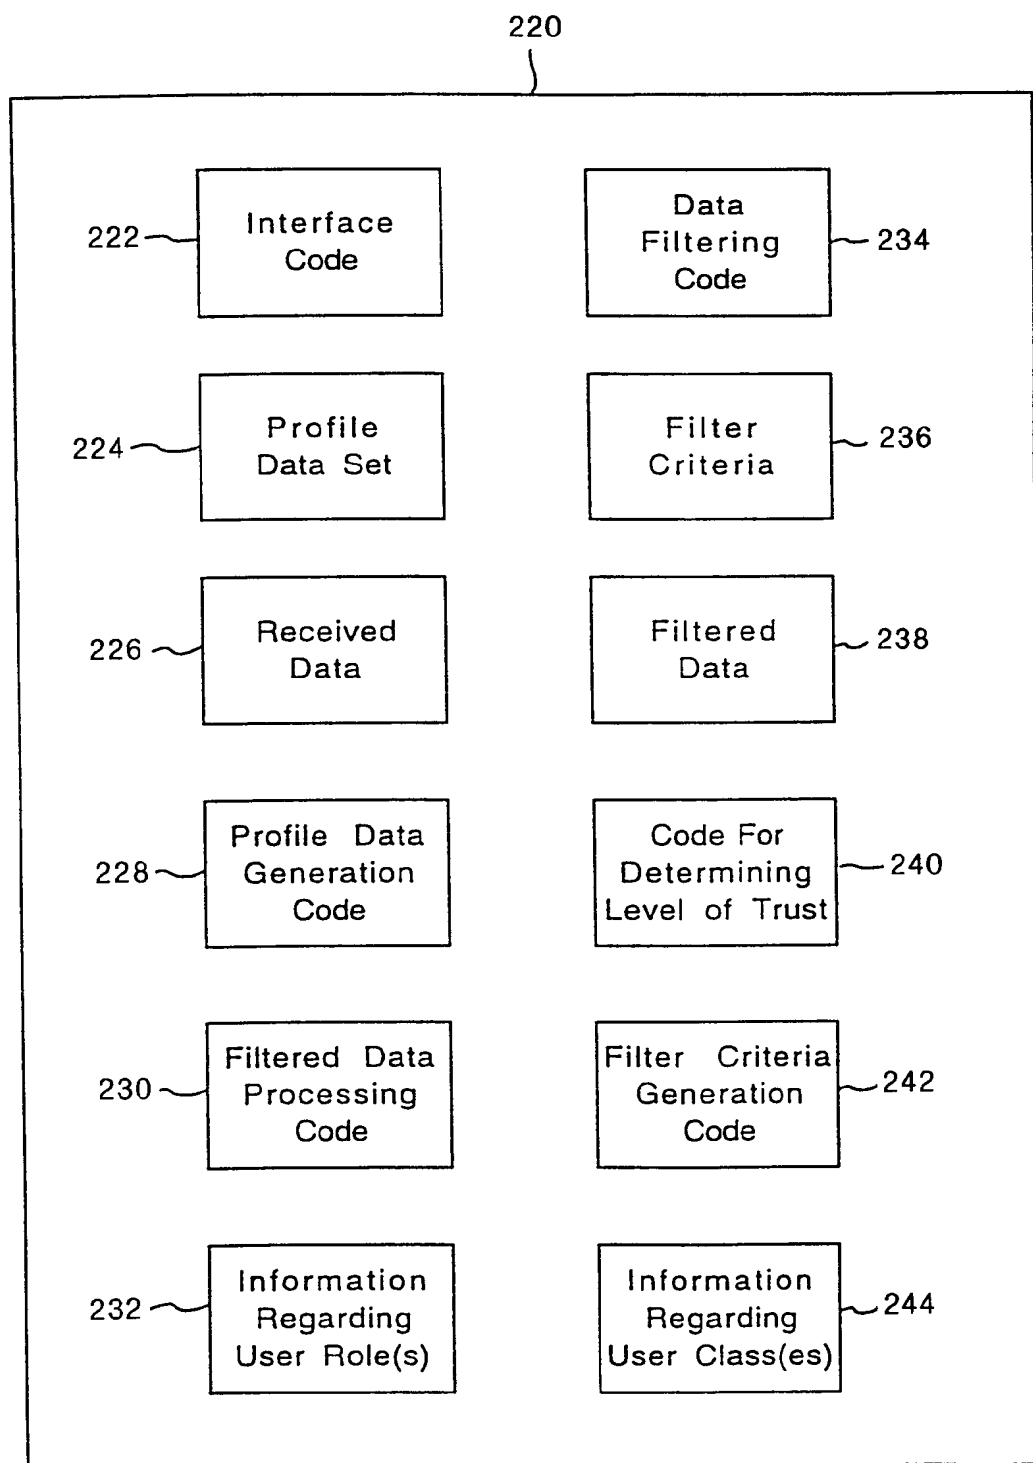


FIG. 15

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/25647

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 17/30, 13/36, 15/00; H04N 1/413

US CL : 395/187/01, 200.61; 707/7, 10

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/187/01, 200.61; 707/7, 10

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS

filter data, filter criteria, profile, client, server

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,606,668 (SHWED) 25 February 97; Abstract; col. 3, lines 44- col. 4, lines 1-43	1-20
Y	US 5,404,505 (LEVINSON) 04 April 95; col. 5, lines 27-55; col. 8, lines 37-46; col. 9, lines 14-25; col. 10, lines 1-35; col. 15, lines 5- 37	1-20
Y,P	US 5,740,423 (LOGAN et al) 14 April 98; abstract; col. 3, lines 49- 64; col. 4, lines 55-66; col. 6, lines 14-61	1-20
A	US 5,596,718 (BOEBERT et al) 21 January 97; abstract; col. 4, lines 7-62	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Z* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

25 JANUARY 1999

Date of mailing of the international search report

15 APR 1999

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 308-5359

Authorized officer

FRANK J. ASTA

Telephone No. (703) 305-3817